USAN: 10/769,079
Filed: January 30, 2004

# BUFFER CONTROL SYSTEM AND METHOD FOR A MEMORY SYSTEM HAVING MEMORY REQUEST BUFFERS

## TECHNICAL FIELD

The present invention is related generally to memory controllers for controlling a memory system, and more particularly, related to memory controllers for managing a memory system having memory request buffers.

## BACKGROUND OF THE INVENTION

Computer systems use memory devices, such as dynamic random access memory ("DRAM") devices, to store data that are accessed by a processor. These memory devices are normally used as system memory in a computer system. In a typical computer system, the processor communicates with the system memory through a processor bus and a memory controller. The memory devices of the system memory, typically arranged in memory modules having multiple memory devices, are coupled through a memory bus to the memory controller. The processor issues a memory request, which includes a memory command, such as a read command, and an address designating the location from which data or instructions are to be read. The memory controller uses the command and address to generate appropriate command signals as well as row and column addresses, which are applied to the system memory through the memory bus. In response to the commands and addresses, data are transferred between the system memory and the processor. The memory controller is often part of a system controller, which also includes bus bridge circuitry for coupling the processor bus to an expansion bus, such as a PCI bus.

In memory systems, high data bandwidth is desirable. Generally, bandwidth limitations are not related to the memory controllers since the memory controllers sequence data to and from the system memory as fast as the memory devices allow. One approach that has been taken to increase bandwidth is to increase the speed of the memory data bus coupling the memory controller to the memory devices. Thus, the same amount of

information can be moved over the memory data bus in less time. However, despite increasing memory data bus speeds, a corresponding increase in bandwidth does not result. One reason for the non-linear relationship between data bus speed and bandwidth is the hardware limitations within the memory devices themselves. That is, the memory controller has to schedule all memory commands to the memory devices such that the hardware limitations are honored. Although these hardware limitations can be reduced to some degree through the design of the memory device, a compromise must be made because reducing the hardware limitations typically adds cost, power, and/or size to the memory devices, all of which are undesirable alternatives. Thus, given these constraints, although it is easy for memory devices to move "well-behaved" traffic at ever increasing rates, for example, sequel traffic to the same page of a memory device, it is much more difficult for the memory devices to resolve "badly-behaved traffic," such as bouncing between different pages or banks of the memory device. As a result, the increase in memory data bus bandwidth does not yield a corresponding increase in information bandwidth.

In addition to the limited bandwidth between processors and memory devices, the performance of computer systems is also limited by latency problems that increase the time required to read data from system memory devices. More specifically, when a memory device read command is coupled to a system memory device, such as a synchronous DRAM ("SDRAM") device, the read data are output from the SDRAM device only after a delay of several clock periods. Therefore, although SDRAM devices can synchronously output burst data at a high data rate, the delay in initially providing the data can significantly slow the operating speed of a computer system using such SDRAM devices. Increasing the memory data bus speed can be used to help alleviate the latency issue. However, as with bandwidth, the increase in memory data bus speeds do not yield a linear reduction of latency, for essentially the same reasons previously discussed.

Although increasing memory data bus speed has, to some degree, been successful in increasing bandwidth and reducing latency, other issues are raised by this

approach. For example, as the speed of the memory data bus increases, loading on the memory bus needs to be decreased in order to maintain signal integrity since traditionally, there has only been wire between the memory controller and the memory slots into which the memory modules are plugged. Several approaches have been taken to accommodate the increase in memory data bus speed. For example, reducing the number of memory slots, adding buffer circuits on a memory module in order to provide sufficient fanout of control signals to the memory devices on the memory module, and providing multiple memory device interfaces on the memory module since there are too few memory module connectors on a single memory device interface. The effectiveness of these conventional approaches are, however, limited. A reason why these techniques were used in the past is that it was cost-effective to do so. However, when only one memory module can be plugged in per interface, it becomes too costly to add a separate memory interface for each required memory slot. In other words, it pushes the system controllers package out of the commodity range and into the boutique range, thereby, greatly adding cost.

One recent approach that allows for increased memory data bus speed in a cost effective manner is the use of multiple memory devices coupled to the processor through a memory hub. In a memory hub architecture, or a hub-based memory sub-system, a system controller or memory controller is coupled over a high speed bi-directional or unidirectional memory controller/hub interface to several memory modules. Typically, the memory modules are coupled in a point-to-point or daisy chain architecture such that the memory modules are connected one to another in series. Thus, the memory controller is coupled to a first memory module, with the first memory module connected to a second memory module, and the second memory module coupled to a third memory module, and so on in a daisy chain fashion.

Each memory module includes a memory hub that is coupled to the memory controller/hub interface and a number of memory devices on the module, with the memory hubs efficiently routing memory requests and responses between the controller and the memory devices over the memory controller/hub interface. Computer systems employing

this architecture can use a high-speed memory data bus since signal integrity can be maintained on the memory data bus. Moreover, this architecture also provides for easy expansion of the system memory without concern for degradation in signal quality as more memory modules are added, such as occurs in conventional memory bus architectures.

Although computer systems using memory hubs may provide superior performance, they may often fail to operate at optimum efficiency for a variety of reasons. One such reason is the issue of managing various buffers in the system memory, for example, memory request buffers and read data buffers included in the memory hubs. Typically, the various buffers are used by the individual memory hubs to handle memory requests in a more efficient manner, such as waiting until a sufficient number of outstanding write requests have been issued to the memory hub before servicing them, or for the purpose of efficiently routing read data sent back to a memory controller from upstream memory hubs after the read requests have been serviced. Monitoring the status of the various buffers of the memory system is desirable to prevent issues such as buffer overflow where continuing to issue memory requests may cause memory requests or read data to stall in the system memory. Additionally, managing the use of the various buffers in the memory system may require accommodating different buffer management goals. For example, it is desirable to allow for a large number of write requests to be issued to the system memory, while minimizing the number of outstanding read requests. Therefore, there is a need for a system and method for managing buffers of a system memory.

## SUMMARY OF THE INVENTION

A memory controller according to one aspect of the present invention includes a memory request queue to receive and store memory requests where the memory request queue issues read requests and write requests to the system memory in response to a flow control signal. The memory controller further includes a response queue coupled to receive a memory request response having a status signal identifying read requests and write requests that have been serviced by the system memory. Further included in the

memory controller is a memory request flow control circuit coupled to the response queue to receive the status signals. The memory request flow control circuit includes separate read and write request monitor circuits to monitor the number of outstanding read and write requests issued to the system memory, respectively. The memory request flow control circuit is also coupled to the memory request queue to provide the flow control signal responsive to the number of outstanding read and write requests to control the issuance of read requests and the issuance of write requests to the system memory.

In another aspect of the invention, a method for managing the issuance of read and write requests to a system memory is provided. The method includes separately monitoring the number of outstanding read requests and write requests issued to the system memory, and separately controlling further issuance of read and write requests to the system memory based on the number of outstanding read and write requests, respectively. In another aspect of the invention, the issuance of read and write requests are managed by independently halting and resuming the issuance of read and write requests to the system memory to maintain the number of outstanding read requests between first and second read thresholds and to maintain the number of outstanding write requests between first and second write thresholds, respectively.

## BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is a partial block diagram of a computer system in which embodiments of the present invention can be implemented.

Figure 2 is a partial block diagram of a memory controller according to an embodiment of the present invention.

Figure 3 is a partial block diagram of a memory controller according to an alternative embodiment of the present invention.

Figure 4 is a partial block diagram of a computer system having a memory hub based system memory in which embodiments of the present invention can be implemented.

## DETAILED DESCRIPTION OF THE INVENTION

Figure 1 illustrates a computer system 100 according to an embodiment of the present invention. The computer system 100 includes a processor 102 for performing various computing functions, such as executing specific software to perform specific calculations or tasks. The processor 102 includes a processor bus 104 that normally includes an address bus, a control bus, and a data bus. The processor bus 104 is typically coupled to a cache memory 106, which, is typically static random access memory ("SRAM") device. The processor bus 104 is further coupled to a system controller 108, which is also sometimes referred to as a bus bridge. The system controller 108 serves as a communications path to the processor 102 for a variety of other components. As shown in Figure 1, the system controller 108 includes a graphics port that is typically coupled to a graphics controller 110, which is, in turn, coupled to a video terminal 112. The system controller 108 is also coupled to one or more input devices 114, such as a keyboard or a mouse, to allow an operator to interface with the computer system 100. Typically, the computer system 100 also includes one or more output devices 116, such as a printer, coupled to the processor 102 through the system controller 108. One or more data storage devices 118 are also typically coupled to the processor 102 through the system controller 108 to allow the processor 102 to store data or retrieve data from internal or external storage media (not shown). Examples of typical storage devices 118 include hard and floppy disks, tape cassettes, and compact disk read-only memories (CD-ROMs).

The system controller 108 includes a memory controller 124 that is coupled to the processor 102. The memory controller 124 is also coupled to a system memory 120 through a system memory bus 122 on which data, address and/or control signals are coupled away from or toward the memory controller 124. The system memory 120 includes a plurality of memory devices 130 coupled to the system memory bus 122, and further includes write buffers 132 and read data buffers 134 also coupled to the system memory bus 122. As with typical computer systems, the system memory 120 is used to store information, typically data or instructions, for use by the rest of the computer system

100, such as, the processor 102, the graphics controller 110, input or output devices 114, 116, and the like. However, the system memory 120 shown in Figure 1 is an "intelligent" memory system that can monitor the types of memory requests issued by the memory controller 124, for example, read requests or write requests, as well as alter the order in which the memory requests are executed from the order in which the memory requests are issued by the memory controller 124. Additionally, the system memory 120 can also modify the flow of read data provided to the memory controller 124 from the order in which the read data is requested by the memory controller 124. The reordering of the memory requests as well as the return read data is made to improve utilization of the system memory bus 122. Handling of the memory requests and read data is facilitated by the use of the write buffer 132 and the read data buffer 134 included in the system memory 120. For example, the write buffer 132 and the read data buffer 134 are used to temporarily store write requests and read data, respectively, in the memory system 120 until a sufficient number of write requests are ready for servicing or until the appropriate time for transferring read data to the memory controller 124.

Further included in the system memory 120 is a memory request response circuit 140 for monitoring the servicing of memory requests issued by the memory request queue 204. In response to the servicing of a memory request, an appropriate response signal is generated by the memory request response circuit 140. For example, when a read request issued to the system memory 120 is serviced, and read data is ready to be provided to the memory controller 124, the memory request response circuit 140 generates a read response signal that includes the read data signals as well as a read status signal that identifies the read request corresponding to the read data. With respect to a write request that has been serviced by the system memory 120, the memory request response circuit 140 generates a write response signal that includes a write status signal indicating that the particular write request has been serviced by the system memory 120. As will be explained in more detail below, the response signals generated by the memory request response circuit are used by the memory controller 124 to monitor the number of outstanding

memory requests issued to the system memory 120. Although in the embodiment described with respect to Figure 1 the system memory 120 includes a memory request response circuit 140 to provide response signals to the memory controller 124, it will be appreciated that alternative approaches known in the art can be used to report the servicing of read and write requests to the memory controller 124. In the interest of brevity, such alternative approaches will not be discussed herein.

An example of a memory system that can be used for the system memory 120 is described in more detail in commonly assigned U.S. Patent Application No. 10/232,473, entitled METHOD AND SYSTEM FOR CONTROLLING MEMORY ACCESSES TO MEMORY MODULES HAVING A MEMORY HUB ARCHITECTURE and filed August 29, 2002 and U.S. Patent Application No. 10/690,810. It will be appreciated by those ordinarily skilled in the art, however, that alternative system memories having buffers for storing memory requests and data can be substituted for the system memory 120 as well, and embodiments of the present invention can be utilized therewith without departing from the scope of the present invention.

Figure 2 illustrates a portion memory controller 202 according to an embodiment of the present invention. The memory controller 202 can be substituted for the memory controller 124 shown in Figure 1. The memory controller 202 includes a memory request queue 204 that receives from the processor 104 (Figure 1) via the system controller 110 high level memory request signals, including command signals, address signals and, in the case of a memory write, write data signals. The memory request queue 204 also receives Request ID signals from a flow control unit 208 that uniquely identify each memory request. These Request IDs are combined with corresponding high level memory requests and stored in the memory request queue 204, preferably in the order they are received. The memory request signals stored in the request queue 204 include both read request signals and write request signals. The high level memory request signals and the Request ID signals will collectively be referred to as memory request signals.

The memory controller 202 also includes a memory response queue 220 that receives read response signals and write response signals from the system memory 120. As previously discussed, the read response signals include read data signals as well as read status signals that identify the read request corresponding to the read data. The write response signals include write status signals that identify a write request that has been serviced by one of the memory modules. The response signals are generated by a memory request response circuit 140 (Figure 1) included in the system memory 120 to allow the memory controller 202 to monitor the number of outstanding read and write requests that have been issued to the system memory 120.

The memory response queue 220 couples read status signals 222 and write status signals 224 to the flow control unit 208 so the flow control unit 208 can determine which read requests and which write requests have been serviced by the system memory 120. The flow control unit 208 makes this determination by comparing the status signals 222, 224 to the Request IDs generated and coupled to the memory request queue 204. As shown in Figure 2, the flow control unit 208 includes read and write buffer control circuits 210, 212, which are used for independently monitoring the number of outstanding read and write requests issued to the system memory 120, respectively. As will be described in more detail below, the flow control unit 208 outputs flow control signals based on the number of outstanding read and write requests to the memory request queue 204 to control whether and when additional read or write requests should be issued to the system memory 120.

In one embodiment of the present invention, the read buffer control circuit 210 and the write buffer control circuit 212 include up/down counter circuits 214, 216 for monitoring outstanding read and write requests issued to the system memory 120. The counter circuit 214 is incremented by the flow control unit 208 in response to the issuance of a read request to the system memory 120, and is decremented in response to the receipt of a read status signal 222 for the corresponding read request. In this manner, the value of the counter circuit 214 is indicative of the number of outstanding read requests issued to

the system memory 120. The counter circuit 216 is used to monitor the number of outstanding write requests in the same manner as previously described with respect to the counter circuit 214 for monitoring the outstanding read requests. Although using the counter circuits 214 and 216 has been described, it will be appreciated, however, that those ordinarily skilled in the art will obtain sufficient understanding from the description provided herein to implement embodiments of the present invention using circuits other than counter circuits 214, 216. Consequently, the scope of the present invention is not limited to the particular embodiment described with respect to Figure 2.

In operation, the flow control unit 208 monitors the number of outstanding read and write requests through the use of the read and write buffer control circuits 210, 212, and the up/down counter circuits 214 and 216. In response to the number of outstanding memory requests, the flow control unit 208 outputs flow control signals to halt or resume further issuance of read or write requests to the system memory 120. In this manner, the flow control unit 208 can prevent overrunning the posted write buffers 132 and the read data buffers 134 of the memory system 120. Additionally, allowing for the number of outstanding read requests before halting further issuance of read requests to be different than the number of outstanding write requests before halting further issuance of write requests provides the flexibility to accommodate differences in buffer management between read and write requests. For example, it is generally the case that a large number of outstanding write requests issued to the system memory 120 is acceptable, and in some cases, preferable, whereas the number of outstanding read requests issued to the system memory 120 is preferably minimized.

Preferably, the read buffer control circuit 210 stores a programmable read request limit that defines the number of outstanding read requests that may be issued to the system memory 120 before the flow control unit 208 outputs a control signal to the memory request queue 204 halting further issuance of read requests. Similarly the write buffer control circuit 212 stores a programmable write request limit that defines the number

of outstanding write requests that may be issued to the system memory 120 before the flow control unit 208 outputs a control signal to halt further issuance of write requests.

During the time issuance of the read or write requests are tolled, outstanding memory requests issued to the system memory 120 are allowed to be serviced, and consequently, clear room in the write buffer and the read data buffer of the system memory 120. With respect to read requests, when the number of outstanding read requests decreases below a programmable read resume limit stored by the read buffer control circuit 210, the flow control unit 208 outputs flow control signals to resume issuance of read requests to the system memory 120. Similarly, with respect to write requests, a programmable write resume limit stored by the write buffer control circuit 212 governs the number of write requests that can be outstanding before resuming the further issuance of write requests to the system memory 120.

Although the read and write request limits, and the read and write resume limit have been previously described as being programmable, it will be appreciated that the read and write request limits and the read and write resume limits can be a fixed value as well. As previously mentioned, programmable read and write buffer limits programmed into the read and write buffer control circuits 210, 212 allow for the buffer control to be tailored for the particular type of memory request, that is, either read requests or write requests, as well as to allow values to be selected that are indicative of the allowable read and write buffer depths of the system memory 120. It will be appreciated that in alternative embodiments, rather than have non-zero read and write resume limits, issuance of read or write requests can resume when the number of outstanding read or write requests decrease to zero.

As previously described with respect to Figure 2, the flow control circuit 208 of the memory controller 202 included read and write buffer control circuits 210, 212 to monitor the number of outstanding read and write requests issued to the system memory 120. In an alternative embodiment of the present invention, shown in Figure 3, the memory controller 302 includes a plurality of read buffer control circuits 310a-d and a

plurality of write buffer control circuits 312a-d. Each of the read buffer control circuits 310a-d has a programmable read request limit and a programmable read resume limit, and each of the write buffer control circuits has a programmable write request limit and a programmable write resume limit. Each of the read buffer control circuits 310a-d can be used to individually monitor a corresponding read data buffer in the system memory 120. Similarly, each of the write buffer control circuits 312a-d can be used to monitor a individually monitor a corresponding write buffer in the system memory 120. In contrast, the memory controller 202 described with respect to Figure 2 monitored the total number of outstanding read requests issued to the system memory 120 using the read buffer control circuit 210 and further monitored the total number of outstanding write requests issued to the system memory 120 using the write buffer control circuit 212.

As with the embodiment of the memory controller 202 described with reference to Figure 2, the individual read and write request limits, and the individual read and write resume limits for the memory controller 302 can be fixed rather than programmable, or alternatively, some combination of fixed and programmable values. Moreover, the read and write buffer control circuits 310a-d, 312a-d can be implemented using counter circuits, as previously described with respect to the memory controller 202. Alternatively, those ordinarily skilled in the art will obtain sufficient understanding from the description provided herein to practice the invention using other well known or later developed designs and circuits. Consequently, the scope of the present invention should not be limited to any of the particular embodiments described herein, or to the use of any particular circuits described with respect to those embodiments.

Figure 4 illustrates a computer system 400 having a memory hub architecture in which embodiments of the present invention can be utilized. The computer system 400 includes a processor 404 for performing various computing functions, such as executing specific software to perform specific calculations or tasks. The processor 404 includes a processor bus 406 that normally includes an address bus, a control bus, and a data bus. The processor bus 406 is typically coupled to cache memory 408, which, is

typically static random access memory ("SRAM"). Finally, the processor bus 406 is coupled to a system controller 410, which is also sometimes referred to as a bus bridge.

The system controller 410 also serves as a communications path to the processor 404 for a variety of other components. More specifically, the system controller 410 includes a graphics port that is typically coupled to a graphics controller 412, which is, in turn, coupled to a video terminal 414. The system controller 410 is also coupled to one or more input devices 418, such as a keyboard or a mouse, to allow an operator to interface with the computer system 400. Typically, the computer system 400 also includes one or more output devices 420, such as a printer, coupled to the processor 404 through the system controller 410. One or more data storage devices 424 are also typically coupled to the processor 404 through the system controller 410 to allow the processor 404 to store data or retrieve data from internal or external storage media (not shown). Examples of typical storage devices 424 include hard and floppy disks, tape cassettes, and compact disk read-only memories (CD-ROMs).

The system controller 410 contains a memory hub controller 428 coupled to several memory modules 430a-n through a bus system 434. Each of the memory modules 430a-n includes a memory hub 440 coupled to several memory devices 448 through command, address and data buses, collectively shown as bus 450. The memory hub 440 efficiently routes memory requests and responses between the controller 428 and the memory devices 448. Each of the memory hubs 440 includes write buffers and read data buffers, as previously described. Computer systems employing this architecture allow for the processor 404 to access one memory module 430a-n while another memory module 430a-n is responding to a prior memory request. For example, the processor 404 can output write data to one of the memory modules 430a-n in the system while another memory module 430a-n in the system is preparing to provide read data to the processor 404. Additionally, a memory hub architecture can also provide greatly increased memory capacity in computer systems.

The memory hub controller 428 includes circuits according to an embodiment of the present invention. That is, outstanding read and write requests are independently monitored, and in response to the number of outstanding read and write requests, the further issuance of read or write requests can be halted until the number of outstanding read or write requests decreases below a resume limit, at which time, issuance of the read or write requests to the memory may resume. In this manner, separate monitoring and control over the issuance of read requests and monitoring and control over the issuance of write requests to the system memory allows for independent buffer control.

From the foregoing it will be appreciated that, although specific embodiments of the invention have been described herein for purposes of illustration, various modifications may be made without deviating from the spirit and scope of the invention. Accordingly, the invention is not limited except as by the appended claims.

# CLAIMS

1. A memory hub controller for controlling access to a system memory, comprising:

a memory request queue storing at least one memory request received through an input port, the memory request queue being operable to issue from an output port each memory request stored in the memory request queue responsive to a flow control signal;

a response queue coupled to receive through an input port read responses containing each read data and a read status signal identifying read requests corresponding to the read data, the response queue further being coupled to receive through the input port write responses each having a write status signal identifying write requests that have been serviced, the response queue being operable to couple at least the read data from each read response signal to a data output port and to couple the read status signal from each read response and the write status signal from each write response to a flow control port; and

a buffer management unit coupled to receive the read status signals and the write status signals from the response queue, the buffer management unit having a read buffer monitor circuit to determine from the read status signals the number of outstanding read requests issued by the memory request queue and having a write buffer monitor circuit to determine from the write status signals the number of outstanding write requests issued by the memory request queue, the buffer management unit further having a flow control circuit coupled to the read and write buffer monitor circuits to generate a flow control signal provided to the memory request queue controlling the issuance of memory requests to the system memory based on the number of outstanding read and write requests issued by the memory request queue.

2. The memory hub controller of claim 1 wherein the flow control circuit of the buffer management unit comprises a control circuit operable to generate a flow control signal halting issuance of read requests to the system memory in response to the read buffer monitor circuit determining the number of outstanding read requests exceeds a read request threshold and

further operable to generate a flow control signal halting issuance of write requests to the system memory in response to the write buffer monitor circuit determining the number of outstanding write requests exceeds a write request threshold.

3. The memory hub controller of claim 2 wherein the read request threshold and the write request threshold are unequal.

4. The memory hub controller of claim 2 wherein the write request threshold is greater than the read request threshold.

5. The memory hub controller of claim 2 wherein the control circuit is further operable to generate a flow control signal resuming issuance of read requests to the system memory in response to the read buffer monitor circuit determining the number of outstanding read requests is less than a read request resume threshold and further operable to generate a flow control signal resuming issuance of write requests to the system memory in response to the write buffer monitor circuit determining the number of outstanding write requests is less than a write request resume threshold.

6. The memory hub controller of claim 2 wherein the control circuit is further operable to generate a flow control signal resuming issuance of read requests to the system memory in response to the read buffer monitor circuit determining the number of outstanding read requests issued to the system memory is zero.

7. The memory hub controller of claim 2 wherein the control circuit is further operable to generate a flow control signal resuming issuance of write requests to the system memory in response to the write buffer monitor circuit determining the number of outstanding write requests issued to the system memory is zero.

8. The memory hub controller of claim 1 wherein the memory request queue is further operable to issue from the output port along with each memory request a request identification uniquely identifying the respective memory request.

9. The memory hub controller of claim 8 wherein the buffer management unit is operable to generate and couple to the memory request queue the request identification for each memory request, and wherein the buffer management unit is further operable to generate the flow control signal on the basis of a comparison between the request identification coupled to the memory request queue and the read status signals and the write status signals received from the response queue.

10. A memory controller coupled to a system memory through a memory bus, comprising:

a memory request queue to receive and store memory requests, the memory request queue issuing read requests and write requests to the system memory in response to a flow control signal;

a response queue coupled to receive a memory request response having a status signal identifying read requests and write requests that have been serviced by the system memory; and

a memory request flow control circuit coupled to the response queue to receive the status signals and having separate read and write request monitor circuits to monitor the number of outstanding read and write requests issued to the system memory, respectively, the memory request flow control circuit coupled to the memory request queue to provide the flow control signal responsive to the number of outstanding read and write requests to control the issuance of read requests and the issuance of write requests to the system memory.

11. The memory controller of claim 10 wherein the memory request flow control circuit comprises a control circuit operable to generate a flow control signal halting

issuance of read requests to the system memory in response to the read request monitor circuit determining the number of outstanding read requests exceeds a read request threshold and further operable to generate a flow control signal halting issuance of write requests to the system memory in response to the write request monitor circuit determining the number of outstanding write requests exceeds a write request threshold.

12. The memory controller of claim 11 wherein the read request threshold and the write request threshold are unequal.

13. The memory controller of claim 11 wherein the write request threshold is greater than the read request threshold.

14. The memory controller of claim 11 wherein the control circuit is further operable to generate a flow control signal resuming issuance of read requests to the system memory in response to the read request monitor circuit determining the number of outstanding read requests is less than a read request resume threshold and further operable to generate a flow control signal resuming issuance of write requests to the system memory in response to the write request monitor circuit determining the number of outstanding write requests is less than a write request resume threshold.

15. The memory controller of claim 11 wherein the control circuit is further operable to generate a flow control signal resuming issuance of read requests to the system memory in response to the read request monitor circuit determining the number of outstanding read requests issued to the system memory is zero.

16. The memory controller of claim 11 wherein the control circuit is further operable to generate a flow control signal resuming issuance of write requests to the system

memory in response to the write request monitor circuit determining the number of outstanding write requests issued to the system memory is zero.

17. The memory controller of claim 10 wherein the memory request queue is further operable to issue along with each memory request a request identification uniquely identifying the respective memory request.

18. The memory controller of claim 17 wherein the memory request flow control circuit is operable to generate and couple to the memory request queue the request identification for each memory request, and wherein the memory request flow control circuit is further operable to generate the flow control signal on the basis of a comparison between the request identification coupled to the memory request queue and the read status signals and the write status signals received from the response queue.

19. A memory hub controller for controlling access to a system memory, comprising:

a memory request queue storing read and write memory requests and being operable to issue to the system memory the read and write memory requests stored in the memory request queue responsive to read and write request flow control signals received at a queue control terminal;

a response queue coupled to receive from the system memory read responses each having a read status signal identifying read requests that have been serviced and to receive write responses each having a write status signal identifying write requests that have been serviced, the response queue being operable to couple the read status signal from each read response and the write status signal from each write response to a flow control port; and

a buffer management unit coupled to the queue control terminal of the memory request queue and further coupled to the flow control port of the response queue, the buffer management unit operable to determine from the read status signals the number of outstanding

read requests issued by the memory request queue and to determine from the write status signals the number of outstanding write requests issued by the memory request queue, and in response to the number of outstanding read and write requests, generate flow control signals to control the issuance of read and write requests by the memory request queue to the system memory.

20. The memory hub controller of claim 19 wherein the buffer management unit comprises a flow control circuit operable to generate a flow control signal halting issuance of read requests to the system memory in response to determining the number of outstanding read requests exceeds a read request threshold and further operable to generate a flow control signal resuming issuance of read requests to the system memory in response to determining the number of outstanding read requests is less than a read request resume threshold.

21. The memory hub controller of claim 19 wherein the buffer management unit comprises a flow control circuit operable to generate a flow control signal halting issuance of write requests to the system memory in response to determining the number of outstanding write requests exceeds a write request threshold and further operable to generate a flow control signal resuming issuance of write requests to the system memory in response to determining the number of outstanding write requests is less than a write request resume threshold.

22. The memory hub controller of claim 19 wherein the memory request queue is further operable to issue along with each memory request a request identification uniquely identifying the respective memory request.

23. The memory hub controller of claim 22 wherein the buffer management unit is operable to generate and couple to the memory request queue the request identification for each memory request, and wherein the buffer management unit is further operable to generate the flow control signal on the basis of a comparison between the request identification coupled to

the memory request queue and the read status signals and the write status signals received from the response queue.

24. A memory system, comprising:

a memory bus for transmitting memory requests and responses thereon;

a plurality of memory modules coupled to the memory bus, each of the modules having:

a plurality of memory devices;

a memory hub coupled to the plurality of devices and the memory bus for receiving memory requests and generating signals in response thereto to operate the plurality of memory devices and service read and write requests, the memory hub including a memory write buffer coupled to the memory bus to store write requests, a memory read queue coupled to the memory devices to receive read data from the memory devices and store the read data for coupling to the memory bus, and further including a response generator to generate a read response including a read status signal identifying read requests corresponding to the read data and to generate a write response including a write status signal identifying write requests that have been serviced; and

a memory controller coupled to the memory bus, the memory controller having:

a memory request queue storing at least one memory request and operable to issue from an output port to the memory modules each memory request stored in the memory request queue responsive to a flow control signal;

a response queue coupled to receive the read responses and write responses, the response queue being operable to couple at least the read data from each read response signal to a data output port and to couple the read status signal from each read response and the write status signal from each write response to a flow control port; and

a buffer management unit coupled to receive the read status signals and the write status signals from the response queue, the buffer management unit having a read buffer monitor circuit to determine from the read status signals the number of outstanding read

requests issued by the memory request queue and having a write buffer monitor circuit to determine from the write status signals the number of outstanding write requests issued by the memory request queue, the buffer management unit further having a flow control circuit coupled to the read and write buffer monitor circuits to generate a flow control signal provided to the memory request queue controlling the issuance of memory requests to the memory modules based on the number of outstanding read and write requests issued by the memory request queue.

25. The memory system of claim 24 wherein the flow control circuit of the buffer management unit comprises a control circuit operable to generate a flow control signal halting issuance of read requests to the system memory in response to the read buffer monitor circuit determining the number of outstanding read requests exceeds a read request threshold and further operable to generate a flow control signal halting issuance of write requests to the system memory in response to the write buffer monitor circuit determining the number of outstanding write requests exceeds a write request threshold.

26. The memory system of claim 25 wherein the read request threshold and the write request threshold are unequal.

27. The memory system of claim 25 wherein the write request threshold is greater than the read request threshold.

28. The memory system of claim 25 wherein the control circuit is further operable to generate a flow control signal resuming issuance of read requests to the system memory in response to the read buffer monitor circuit determining the number of outstanding read requests is less than a read request resume threshold and further operable to generate a flow control signal resuming issuance of write requests to the system memory in response to the write buffer monitor circuit determining the number of outstanding write requests is less than a write request resume threshold.

29. The memory system of claim 25 wherein the control circuit is further operable to generate a flow control signal resuming issuance of read requests to the system memory in response to the read buffer monitor circuit determining the number of outstanding read requests issued to the system memory is zero.

30. The memory system of claim 25 wherein the control circuit is further operable to generate a flow control signal resuming issuance of write requests to the system memory in response to the write buffer monitor circuit determining the number of outstanding write requests issued to the system memory is zero.

31. The memory system of claim 24 wherein the memory request queue is further operable to issue from the output port along with each memory request a request identification uniquely identifying the respective memory request.

32. The memory system of claim 31 wherein the buffer management unit is operable to generate and couple to the memory request queue the request identification for each memory request, and wherein the buffer management unit is further operable to generate the flow control signal on the basis of a comparison between the request identification coupled to the memory request queue and the read status signals and the write status signals received from the response queue.

33. A memory system, comprising:

a memory bus for transmitting memory requests and responses thereon;

a system memory having a plurality of memory modules coupled to the memory bus, each of the modules having:

a plurality of memory devices;

a memory hub coupled to the plurality of devices and the memory bus for receiving memory requests and generating signals in response thereto to operate the plurality of

memory devices and service read and write requests, the memory hub including a memory write buffer coupled to the memory bus to store write requests, a memory read queue coupled to the memory devices to receive read data from the memory devices and store the read data for coupling to the memory bus, and further including a response generator to generate a read response including a read status signal identifying read requests corresponding to the read data and to generate a write response including a write status signal identifying write requests that have been serviced; and

a memory controller coupled to the system memory through the memory bus, the memory controller having:

a memory request queue to receive and store memory requests, the memory request queue issuing read requests and write requests to the system memory in response to a flow control signal;

a response queue coupled to receive a memory request response having a status signal identifying read requests and write requests that have been serviced by the system memory; and

a memory request flow control circuit coupled to the response queue to receive the status signals and having separate read and write request monitor circuits to monitor the number of outstanding read and write requests issued to the system memory, respectively, the memory request flow control circuit coupled to the memory request queue to provide the flow control signal responsive to the number of outstanding read and write requests to control the issuance of read requests and the issuance of write requests to the system memory.

34. The memory system of claim 33 wherein the memory request flow control circuit comprises a control circuit operable to generate a flow control signal halting issuance of read requests to the system memory in response to the read request monitor circuit determining the number of outstanding read requests exceeds a read request threshold and further operable to generate a flow control signal halting issuance of write requests to the system memory in

response to the write request monitor circuit determining the number of outstanding write requests exceeds a write request threshold.

35. The memory system of claim 34 wherein the read request threshold and the write request threshold are unequal.

36. The memory system of claim 34 wherein the write request threshold is greater than the read request threshold.

37. The memory system of claim 34 wherein the control circuit is further operable to generate a flow control signal resuming issuance of read requests to the system memory in response to the read request monitor circuit determining the number of outstanding read requests is less than a read request resume threshold and further operable to generate a flow control signal resuming issuance of write requests to the system memory in response to the write request monitor circuit determining the number of outstanding write requests is less than a write request resume threshold.

38. The memory system of claim 34 wherein the control circuit is further operable to generate a flow control signal resuming issuance of read requests to the system memory in response to the read request monitor circuit determining the number of outstanding read requests issued to the system memory is zero.

39. The memory system of claim 34 wherein the control circuit is further operable to generate a flow control signal resuming issuance of write requests to the system memory in response to the write request monitor circuit determining the number of outstanding write requests issued to the system memory is zero.

40. The memory system of claim 33 wherein the memory request queue is further operable to issue along with each memory request a request identification uniquely identifying the respective memory request.

41. The memory system of claim 40 wherein the memory request flow control circuit is operable to generate and couple to the memory request queue the request identification for each memory request, and wherein the memory request flow control circuit is further operable to generate the flow control signal on the basis of a comparison between the request identification coupled to the memory request queue and the read status signals and the write status signals received from the response queue.

42. A processor-based system, comprising:

a processor having a processor bus;

a system controller coupled to the processor bus, the system controller having a peripheral device port, the system controller further comprising a controller coupled to a system memory port;

at least one input device coupled to the peripheral device port of the system controller;

at least one output device coupled to the peripheral device port of the system controller;

at least one data storage device coupled to the peripheral device port of the system controller; and

a memory system, comprising:

a memory bus for transmitting memory requests and responses thereon;

a plurality of memory modules coupled to the memory bus, each of the modules having:

a plurality of memory devices;

a memory hub coupled to the plurality of devices and the memory bus for receiving memory requests and generating signals in response thereto to operate the plurality of memory devices and service read and write requests, the memory hub including a memory write buffer coupled to the memory bus to store write requests, a memory read queue coupled to the memory devices to receive read data from the memory devices and store the read data for coupling to the memory bus, and further including a response generator to generate a read response including a read status signal identifying read requests corresponding to the read data and to generate a write response including a write status signal identifying write requests that have been serviced; and

a memory controller coupled to the memory bus, the memory controller having:

a memory request queue storing at least one memory request and operable to issue from an output port to the memory modules each memory request stored in the memory request queue responsive to a flow control signal;

a response queue coupled to receive the read responses and write responses, the response queue being operable to couple at least the read data from each read response signal to a data output port and to couple the read status signal from each read response and the write status signal from each write response to a flow control port; and

a buffer management unit coupled to receive the read status signals and the write status signals from the response queue, the buffer management unit having a read buffer monitor circuit to determine from the read status signals the number of outstanding read requests issued by the memory request queue and having a write buffer monitor circuit to determine from the write status signals the number of outstanding write requests issued by the memory request queue, the buffer management unit further having a flow control circuit coupled to the read and write buffer monitor circuits to generate a flow control signal provided to the memory request queue controlling the issuance of memory requests to the memory modules based on the number of outstanding read and write requests issued by the memory request queue.

43. The processor-based system of claim 42 wherein the flow control circuit of the buffer management unit comprises a control circuit operable to generate a flow control signal halting issuance of read requests to the system memory in response to the read buffer monitor circuit determining the number of outstanding read requests exceeds a read request threshold and further operable to generate a flow control signal halting issuance of write requests to the system memory in response to the write buffer monitor circuit determining the number of outstanding write requests exceeds a write request threshold.

44. The processor-based system of claim 43 wherein the read request threshold and the write request threshold are unequal.

45. The processor-based system of claim 43 wherein the write request threshold is greater than the read request threshold.

46. The processor-based system of claim 43 wherein the control circuit is further operable to generate a flow control signal resuming issuance of read requests to the system memory in response to the read buffer monitor circuit determining the number of outstanding read requests is less than a read request resume threshold and further operable to generate a flow control signal resuming issuance of write requests to the system memory in response to the write buffer monitor circuit determining the number of outstanding write requests is less than a write request resume threshold.

47. The processor-based system of claim 43 wherein the control circuit is further operable to generate a flow control signal resuming issuance of read requests to the system memory in response to the read buffer monitor circuit determining the number of outstanding read requests issued to the system memory is zero.

48. The processor-based system of claim 43 wherein the control circuit is further operable to generate a flow control signal resuming issuance of write requests to the system memory in response to the write buffer monitor circuit determining the number of outstanding write requests issued to the system memory is zero.

49. The processor-based system of claim 42 wherein the memory request queue is further operable to issue from the output port along with each memory request a request identification uniquely identifying the respective memory request.

50. The processor-based system of claim 49 wherein the buffer management unit is operable to generate and couple to the memory request queue the request identification for each memory request, and wherein the buffer management unit is further operable to generate the flow control signal on the basis of a comparison between the request identification coupled to the memory request queue and the read status signals and the write status signals received from the response queue.

51. A processor-based system, comprising:

a processor having a processor bus;

a system controller coupled to the processor bus, the system controller having a peripheral device port, the system controller further comprising a controller coupled to a system memory port;

at least one input device coupled to the peripheral device port of the system controller;

at least one output device coupled to the peripheral device port of the system controller;

at least one data storage device coupled to the peripheral device port of the system controller; and

a memory system, comprising:

a memory bus for transmitting memory requests and responses thereon;

a system memory having a plurality of memory modules coupled to the memory bus, each of the modules having:

a plurality of memory devices;

a memory hub coupled to the plurality of devices and the memory bus for receiving memory requests and generating signals in response thereto to operate the plurality of memory devices and service read and write requests, the memory hub including a memory write buffer coupled to the memory bus to store write requests, a memory read queue coupled to the memory devices to receive read data from the memory devices and store the read data for coupling to the memory bus, and further including a response generator to generate a read response including a read status signal identifying read requests corresponding to the read data and to generate a write response including a write status signal identifying write requests that have been serviced; and

a memory controller coupled to the system memory through the memory bus, the memory controller having:

a memory request queue to receive and store memory requests, the memory request queue issuing read requests and write requests to the system memory in response to a flow control signal;

a response queue coupled to receive a memory request response having a status signal identifying read requests and write requests that have been serviced by the system memory; and

a memory request flow control circuit coupled to the response queue to receive the status signals and having separate read and write request monitor circuits to monitor the number of outstanding read and write requests issued to the system memory, respectively, the memory request flow control circuit coupled to the memory request queue to provide the flow control signal responsive to the number of outstanding read and write requests to control the issuance of read requests and the issuance of write requests to the system memory.

52. The processor-based system of claim 51 wherein the memory request flow control circuit comprises a control circuit operable to generate a flow control signal halting issuance of read requests to the system memory in response to the read request monitor circuit determining the number of outstanding read requests exceeds a read request threshold and further operable to generate a flow control signal halting issuance of write requests to the system memory in response to the write request monitor circuit determining the number of outstanding write requests exceeds a write request threshold.

53. The processor-based system of claim 52 wherein the read request threshold and the write request threshold are unequal.

54. The processor-based system of claim 52 wherein the write request threshold is greater than the read request threshold.

55. The processor-based system of claim 52 wherein the control circuit is further operable to generate a flow control signal resuming issuance of read requests to the system memory in response to the read request monitor circuit determining the number of outstanding read requests is less than a read request resume threshold and further operable to generate a flow control signal resuming issuance of write requests to the system memory in response to the write request monitor circuit determining the number of outstanding write requests is less than a write request resume threshold.

56. The processor-based system of claim 52 wherein the control circuit is further operable to generate a flow control signal resuming issuance of read requests to the system memory in response to the read request monitor circuit determining the number of outstanding read requests issued to the system memory is zero.

57. The processor-based system of claim 52 wherein the control circuit is further operable to generate a flow control signal resuming issuance of write requests to the system memory in response to the write request monitor circuit determining the number of outstanding write requests issued to the system memory is zero.

58. The processor-based system of claim 51 wherein the memory request queue is further operable to issue along with each memory request a request identification uniquely identifying the respective memory request.

59. The processor-based system of claim 58 wherein the memory request flow control circuit is operable to generate and couple to the memory request queue the request identification for each memory request, and wherein the memory request flow control circuit is further operable to generate the flow control signal on the basis of a comparison between the request identification coupled to the memory request queue and the read status signals and the write status signals received from the response queue.

60. A method for managing the issuance of read and write requests to a system memory, comprising:

separately monitoring the number of outstanding read requests and write requests issued to the system memory;

separately controlling further issuance of read and write requests to the system memory based on the number of outstanding read and write requests, respectively.

61. The method of claim 60 wherein separately controlling further issuance of read and write requests to the system memory based on the number of outstanding read and write requests comprises halting further issuance of read requests in response to the number of outstanding read requests exceeds a read request threshold and halting further issuance of write

requests in response to the number of outstanding write requests exceeds a write request threshold.

62. The method of claim 61 wherein the read request threshold and the write request threshold are unequal.

63. The method of claim 62 wherein the write request threshold is greater than the read request threshold.

64. The method of claim 61, further comprising resuming issuance of read requests in response to the number of outstanding read requests decreasing to less than a read resume threshold.

65. The method of claim 61, further comprising resuming issuance of write request in response to the number of outstanding write requests decreasing to less than a write resume threshold.

66. The method of claim 60 wherein separately monitoring the number of outstanding read requests and write requests issued to the system memory comprises incrementing a read request counter in response to issuance of a read request and decrementing the read request counter in response to receiving an indication that the corresponding read request has been serviced by the system memory and incrementing a write request counter in response to issuance of a write request and decrementing the write request counter in response to receiving an indication that the corresponding write request has been serviced by the system memory.

67. A method for managing the issuance of read and write requests to a system memory, comprising:

monitoring the number of outstanding read requests and write requests issued to the system memory; and

independently halting and resuming the issuance of read and write requests to the system memory to maintain the number of outstanding read requests between first and second read thresholds and to maintain the number of outstanding write requests between first and second write thresholds, respectively.

68. The method of claim 67 wherein the first read threshold is greater than the second read threshold and the first write threshold is greater than the second write threshold, and the first write threshold is unequal to the first read threshold.

69. The method of claim 67 wherein resuming issuance of read requests to the system memory comprises resuming issuance of the read requests in response to the number of outstanding read requests being equal to zero.

70. The method of claim 67 wherein resuming issuance of write requests to the system memory comprises resuming issuance of the write requests in response to the number of outstanding write requests being equal to zero.

71. The method of claim 67, further comprising generating for each memory request a request identification uniquely identifying the respective memory request.

72. The method of claim 71 wherein monitoring the number of outstanding read and write requests comprises monitoring the servicing of the particular read and write requests issued to the system memory based on the request identification.

73. The method of claim 67 wherein monitoring the number of outstanding read requests and write requests issued to the system memory comprises incrementing a read request counter in response to issuance of a read request and decrementing the read request counter in response to receiving an indication that the corresponding read request has been serviced by the system memory and incrementing a write request counter in response to issuance of a write request and decrementing the write request counter in response to receiving an indication that the corresponding write request has been serviced by the system memory.

74. A method for managing the issuance of read and write requests to a system memory, comprising:

monitoring the number of outstanding read requests issued to the system memory;

halting the issuance of read requests to the system memory in response to the number of outstanding read requests exceeding a read request threshold;

resuming the issuance of read requests to the system memory in response to the number of outstanding read requests decreasing to less than a read resume threshold;

monitoring the number of outstanding write requests issued to the system memory;

halting the issuance of write requests to the system memory in response to the number of outstanding write requests exceeding a write request threshold; and

resuming the issuance of write requests to the system memory in response to the number of outstanding write requests decreasing to less than a write resume threshold.

75. The method of claim 74 wherein the read request threshold and the write request threshold are unequal.

76. The method of claim 75 wherein the write request threshold is greater than the read request threshold.

77. The method of claim 74 wherein resuming issuance of read requests to the system memory comprises resuming issuance of the read requests in response to the number of outstanding read requests being equal to zero.

78. The method of claim 74 wherein resuming issuance of write requests to the system memory comprises resuming issuance of the write requests in response to the number of outstanding write requests being equal to zero.

79. The method of claim 74, further comprising generating for each read and write request a request identification uniquely identifying the respective memory request.

80. The method of claim 79 wherein monitoring the number of outstanding read and write requests comprises monitoring the servicing of the particular read and write requests issued to the system memory based on the request identification.

81. The method of claim 74 wherein monitoring the number of outstanding read requests issued to the system memory comprises incrementing a read request counter in response to issuance of a read request and decrementing the read request counter in response to receiving an indication that the corresponding read request has been serviced by the system memory and monitoring the number of outstanding write requests issued to the system memory comprises incrementing a write request counter in response to issuance of a write request and decrementing the write request counter in response to receiving an indication that the corresponding write request has been serviced by the system memory.

# BUFFER CONTROL SYSTEM AND METHOD FOR A MEMORY SYSTEM HAVING MEMORY REQUEST BUFFERS

## ABSTRACT OF THE DISCLOSURE

A memory controller and method for managing the issuance of read and write requests to a system memory is provided. The number of outstanding read requests and write requests issued to the system memory are separately monitored and further issuance of read and write requests to the system memory are separately controlled based on the number of outstanding read and write requests, respectively. For example, the issuance of read and write requests can be managed by independently halting and resuming the issuance of read and write requests to the system memory to maintain the number of outstanding read requests between first and second read thresholds and to maintain the number of outstanding write requests between first and second write thresholds, respectively.

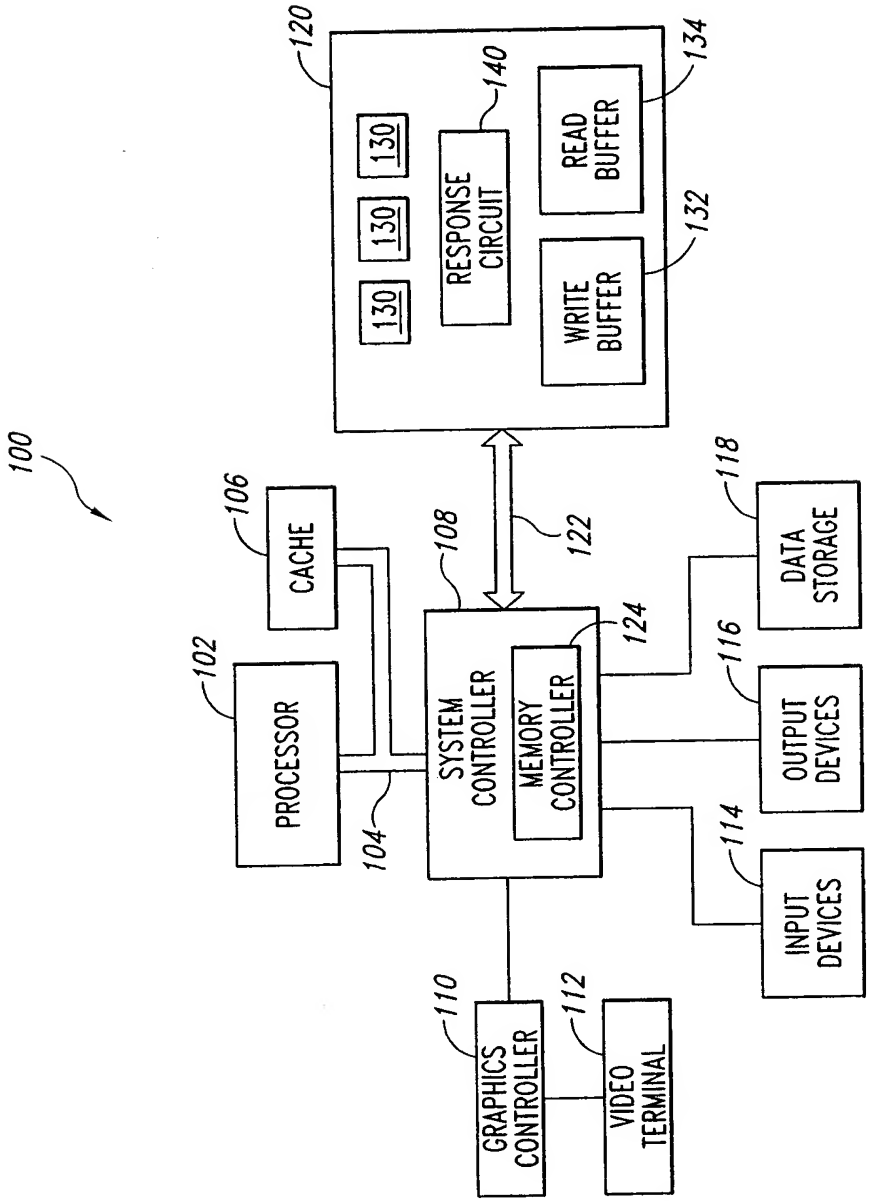
100
120
130
130
130
RESPONSE CIRCUIT
140
READ BUFFER
134
WRITE BUFFER
132
106
CACHE
108
122
102
PROCESSOR
104
SYSTEM CONTROLLER
MEMORY CONTROLLER
124
118
DATA STORAGE
116
OUTPUT DEVICES
114
INPUT DEVICES
110
GRAPHICS CONTROLLER
112
VIDEO TERMINAL

*Fig. 1*

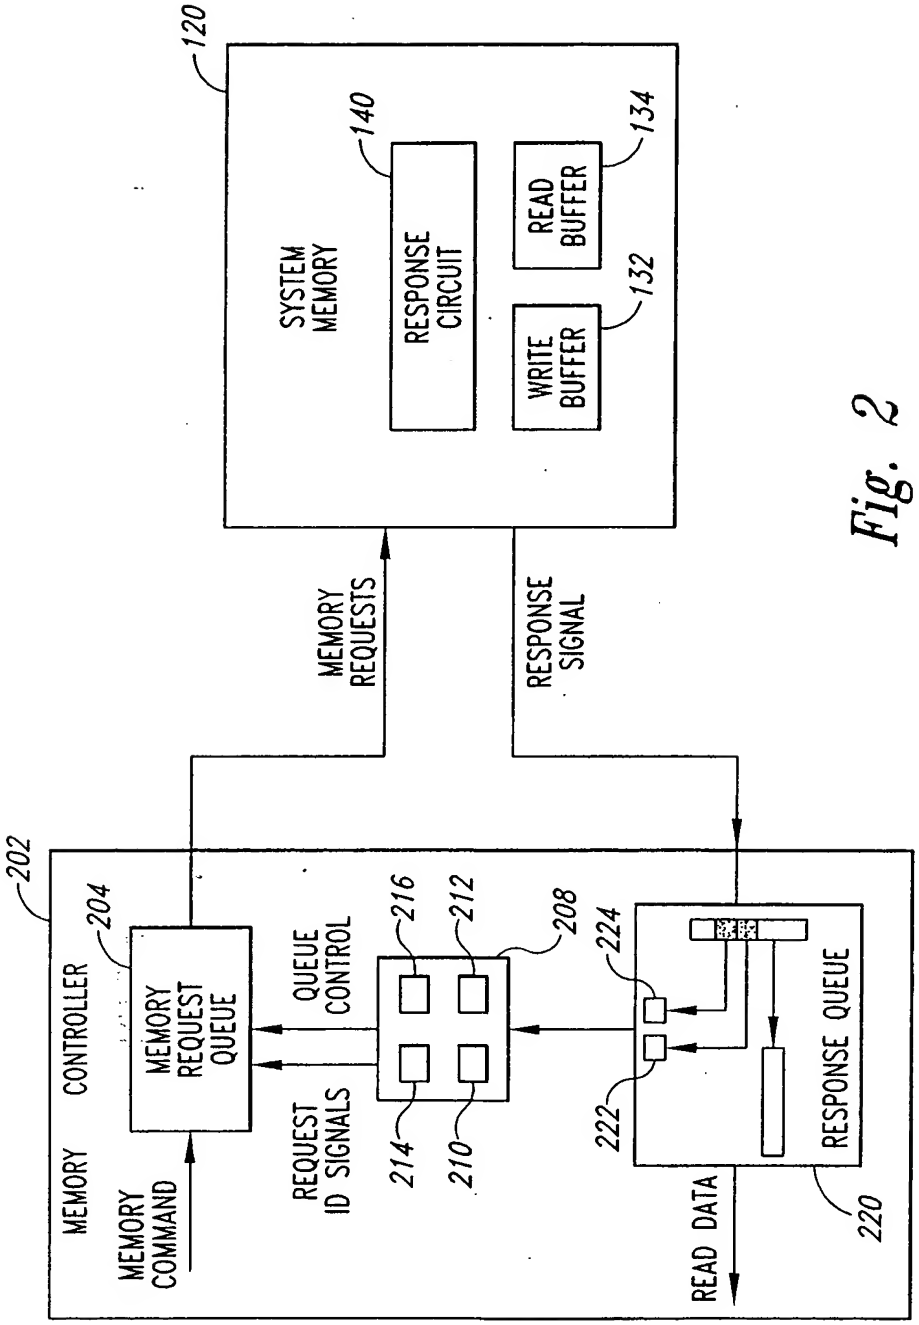
120
SYSTEM MEMORY
140
RESPONSE CIRCUIT
134
READ BUFFER
132
WRITE BUFFER
MEMORY REQUESTS
RESPONSE SIGNAL
202
MEMORY CONTROLLER
204
MEMORY REQUEST QUEUE
QUEUE CONTROL
216
212
208
224
REQUEST ID SIGNALS
214
210
222
RESPONSE QUEUE
MEMORY COMMAND
READ DATA
220

Fig. 2

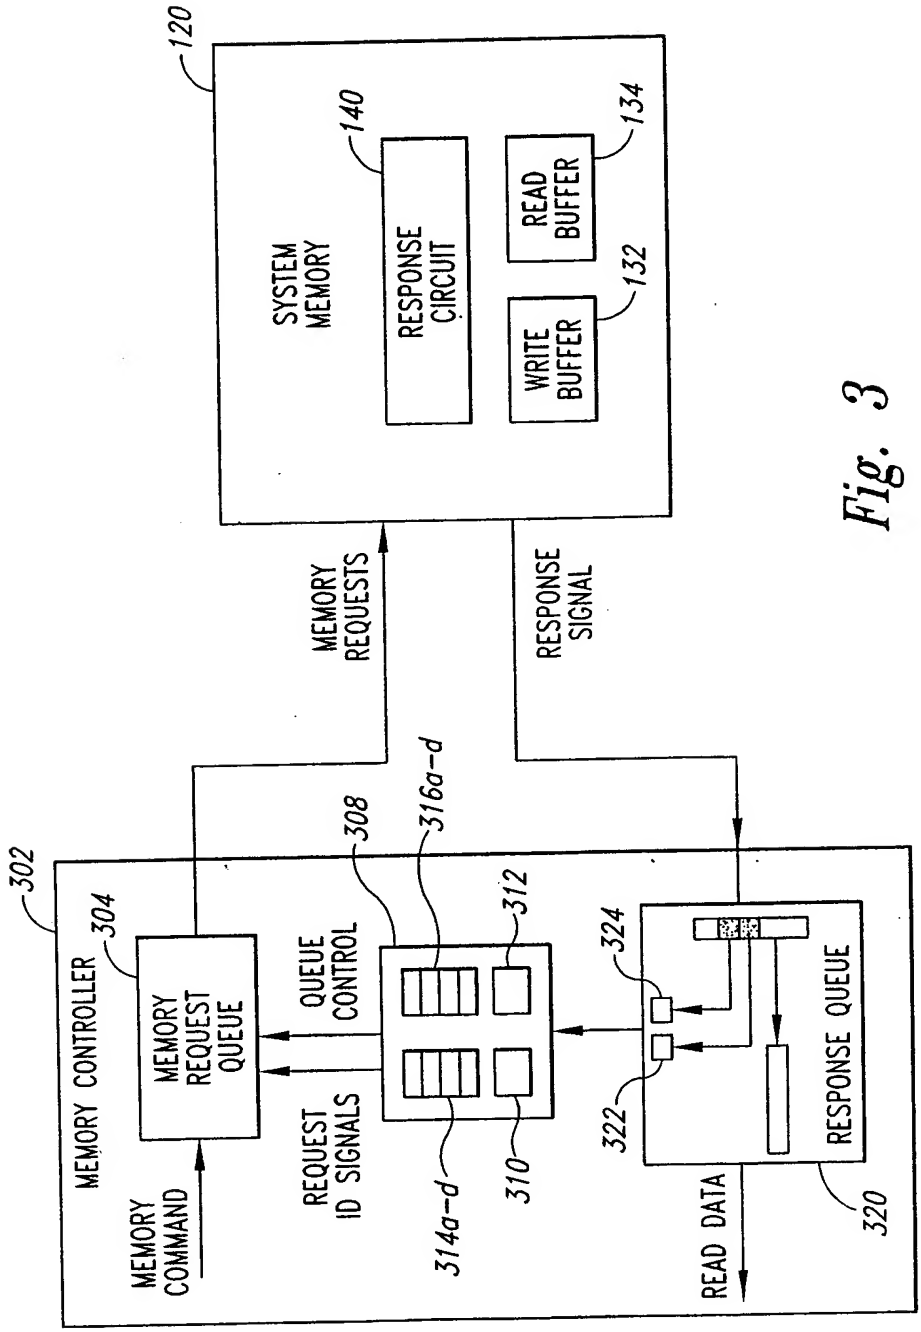
120
SYSTEM MEMORY
140
RESPONSE CIRCUIT
134
READ BUFFER
132
WRITE BUFFER
MEMORY REQUESTS
RESPONSE SIGNAL
302
MEMORY CONTROLLER
304
MEMORY REQUEST QUEUE
QUEUE CONTROL
308
316a-d
312
324
MEMORY COMMAND
REQUEST ID SIGNALS
314a-d
310
322
RESPONSE QUEUE
READ DATA
320

Fig. 3

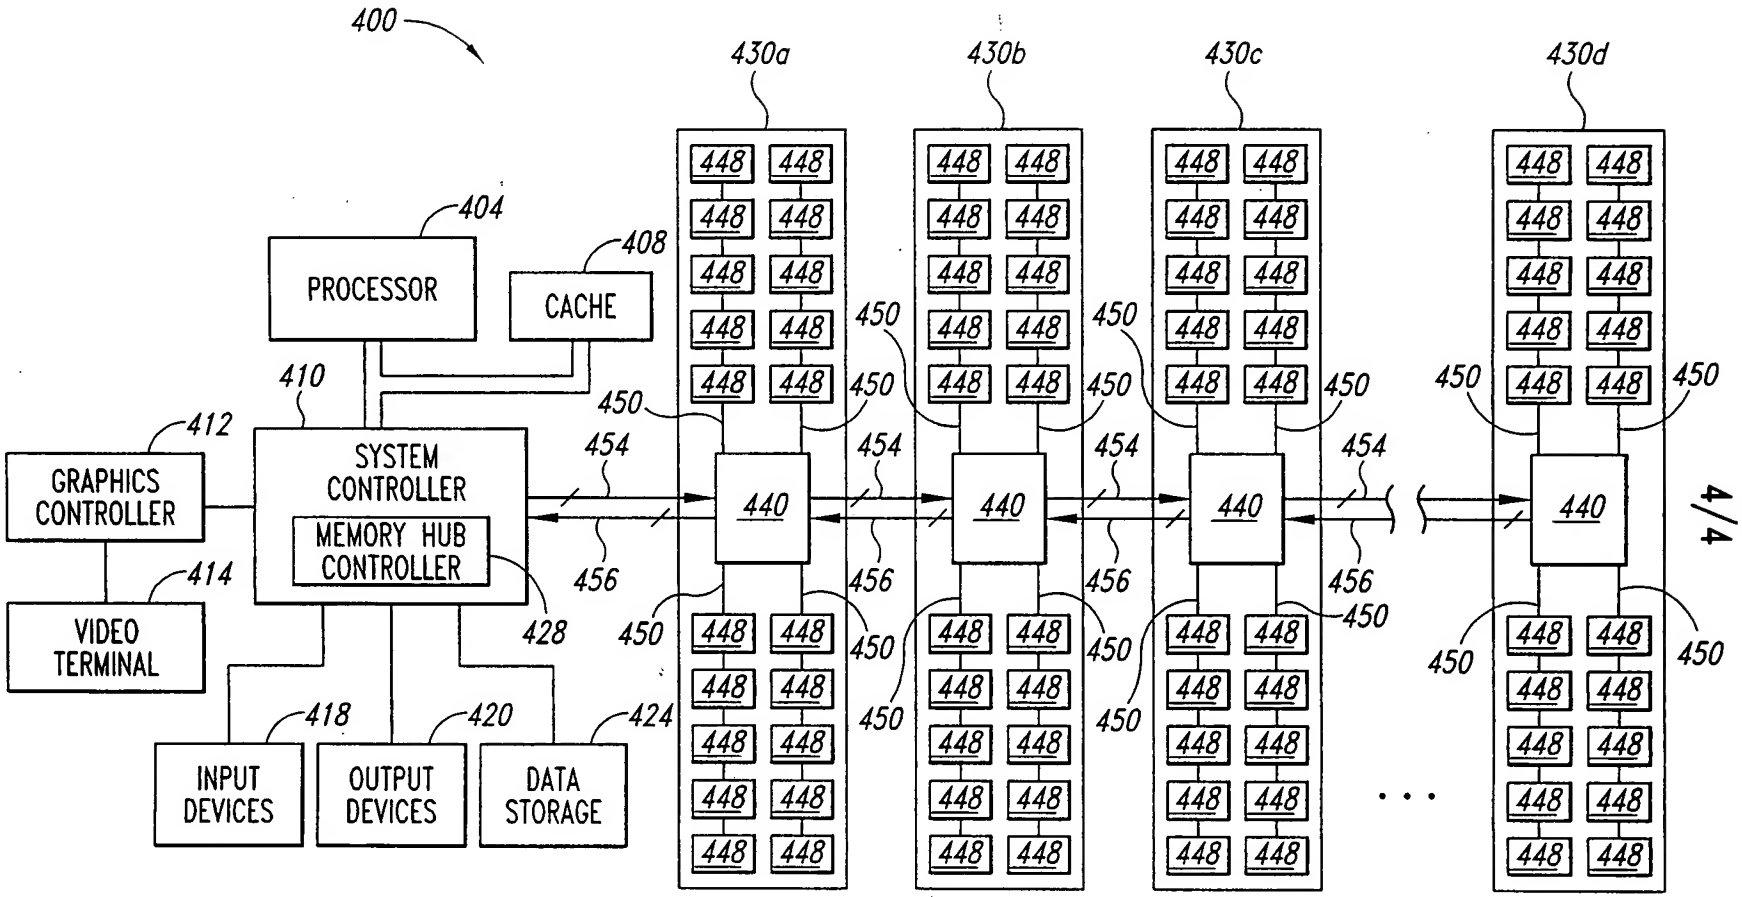
400
430a
430b
430c
430d
404
PROCESSOR
408
CACHE
410
412
GRAPHICS CONTROLLER
SYSTEM CONTROLLER
MEMORY HUB CONTROLLER
428
414
VIDEO TERMINAL
418
INPUT DEVICES
420
OUTPUT DEVICES
424
DATA STORAGE
440
448
450
454
456
4/4

*Fig. 4*